Aspects of Sharing Flight Data via SWIM

Integrated CNS Conference & Workshop

May, 2005

Jon Dehn
Lockheed Martin Transportation & Security Solutions
Jon.Dehn@lmco.com

Sid Rudolph
IT Consulting Group
Sid.Rudolph@lmco.com

LOCKHEED MARTIN

Overview

- Each of the major data types proposed for SWIM has unique aspects
- Discussed here are some of those aspects for Flight Data
- Techniques for handling Flight Data may have implications for the SWIM infrastructure, and may be of interest for other data types

Flight Data in Context

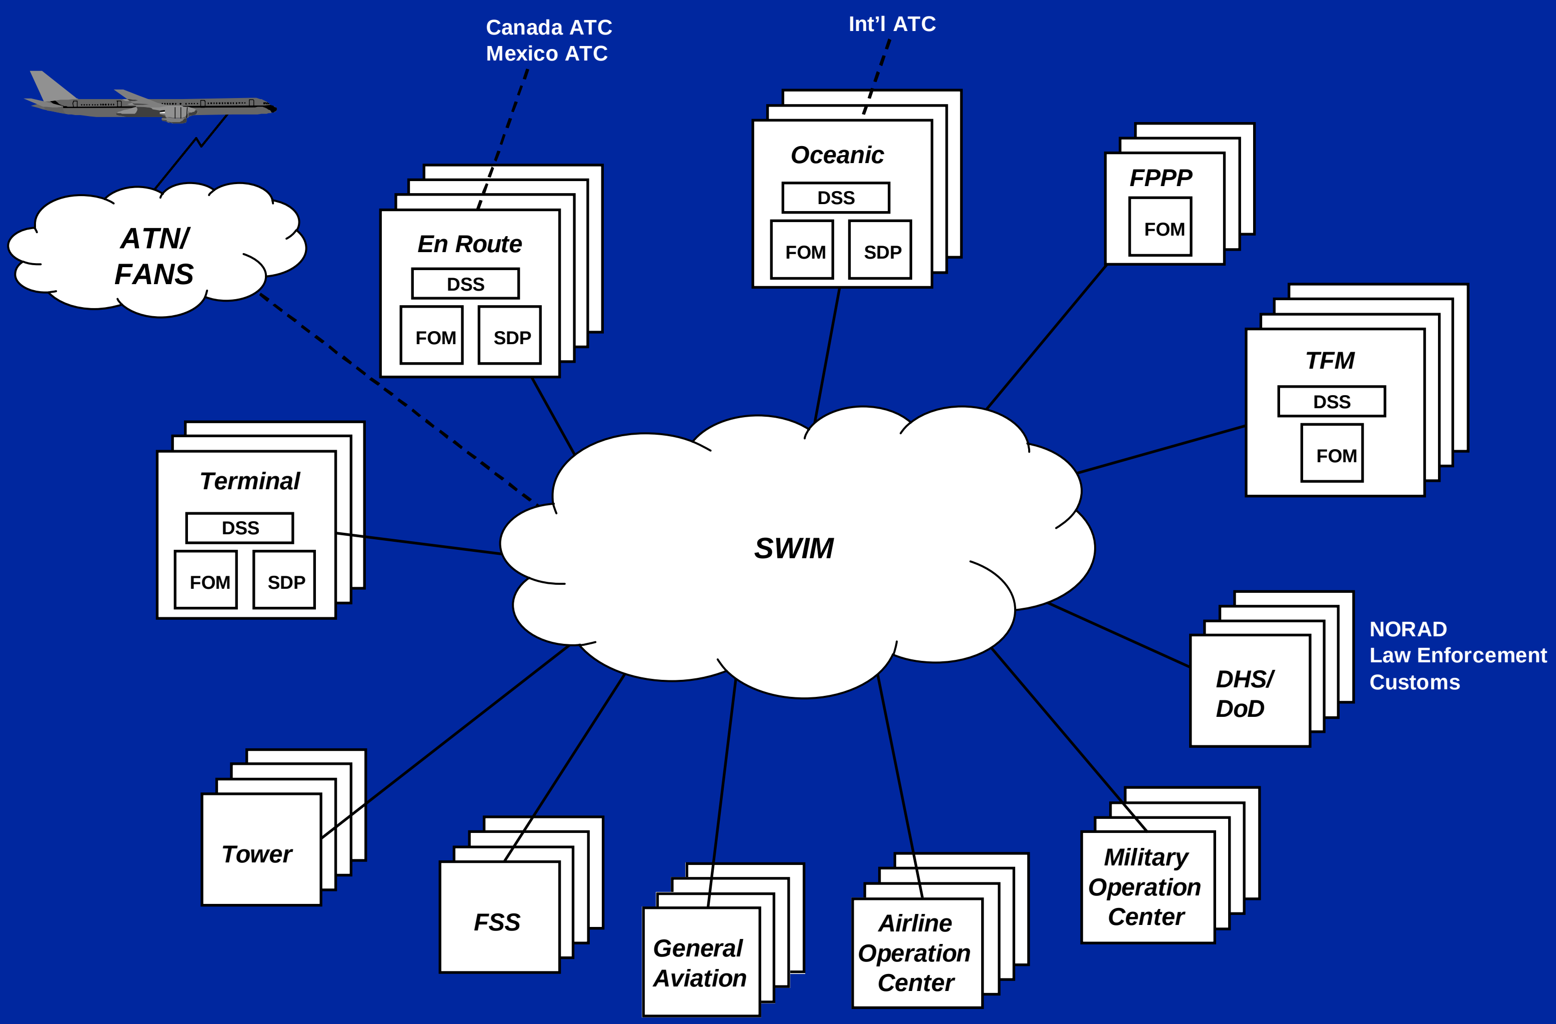

Flight Data Topics

- Diverse Sources of Information
- Volume of Data
- Client Initialization
- Security
- Relationships to Other NAS Information
- Extensibility

Diverse Sources of Information

- NAS systems pass "ownership" of a flight from installation to installation
- This can be extended in the pre-flight phase to the flow management / strategic planning systems
- In order to manage updates to the data and to provide a single, authoritative copy of flight data, there must be a mediation process to control when the different systems can update the various fields of a flight (SWIM clients shouldn't care about NAS system ownership of data)

Diverse Sources (2 of 2)

- The aircraft itself will be the source for some data
- Presumably, for this subset, the aircraft will be considered to be the authoritative source, regardless of which ground system happens to be in control of the flight at the time
- A natural design would be a SWIM-Enabled application that:
  - received updates (including ownership changes) from various ground and airborne systems
  - Applied the updates according to the ownership rules, and
  - Published the resulting data to SWIM

Volume of Data

- Different SWIM data types have different characteristics:
  - Surveillance has a high volume of small updates
  - Aeronautical data has a low volume of large updates
  - Flight data falls somewhere in between

Volume of Data (2 of 3)

- Considering all data for a flight, including one or more 4-D trajectories, a single set of flight information is non-trivial in size
  - In binary form, it could easily exceed 100K bytes
  - In XML form, it would be considerably bigger
- Different data fields get updated at different rates
  - The position of the flight changes rapidly
  - The filed route changes infrequently

Volume of Data (3 of 3)

- The raw size and variation in update rate will cause us to consider:
  - Selective subscription by field:
    - clients only get the subset of the flight's fields they are interested in
    - Only get an update when one of those fields changes
  - Selective subscription by flight:
    - Clients can select the subset of flights according to geography, time, airline, or other flight attributes
  - Client specification of the update rate (perhaps the current position is needed only every minute, not every time it changes)

Client Initialization

- For surveillance data, which constantly refreshes itself, it may be sufficient just to listen in to the current data stream to get a full system picture
- For flight data, that's not the case. It must be possible to get the current set of flights during client initialization
- This raises several design issues

Client Initialization (2 of 2)

- Some of the issues:
  - Server support of steady state updates along with one or more client initialization requests
  - Network load considerations during the initialization phase
  - Client synchronization of updates with initialization data received during the (probably) lengthy initialization process
- Assumptions must be made about:
  - Number of clients
  - Network bandwidth available
  - Server processing power available

Security of Data

- All SWIM supplied data will have general security restrictions, such as not supplying DOD or HLS sensitive data to the general public
- Flight data contains business sensitive data, which must also be protected
- It may be necessary to sanitize the data so that, for example, an AOC could see detailed information about their own flights, but only generic information about other flights. This allows the AOC to make decisions based on the general traffic picture, but not know too much about other airlines operations

Relationships to Other NAS Information

- Parts of flight data are derived from different lower level data
- Clients may want to know what “version” of the lower level data was used
- Examples:
  - The route of flight is calculated based on a set of preferential routes. The active preferential routes can change dynamically
  - The trajectory is modeled based on wind data. That is updated periodically (once/hour).

Extensibility

- All SWIM data must be extensible; flight data will probably have the highest rate of additions and changes to the underlying data structures
- It's likely SWIM flight data will be described by a set of XML Schemas
  - Simple additions to the schema can be made without impact to clients if the clients have been designed to ignore unrecognized fields
  - More complex changes will require schema version management

Extensibility (2 of 2)

- The publisher of SWIM Flight Data must support some number of schema versions concurrently
  – Clients would specify the desired version when they subscribed to flight data
- As with all SWIM data, there must be some "control board" to decide what each version of the schema will contain, and how many back-level versions will be supported

Summary

- Many of these same issues are faced in existing system designs (such as ERAM, STARS, etc), but those programs have a more controlled design process
  - Consensus is more easily reached on approaches
  - Resources and requirements are more tightly controlled
- The order-of-magnitude increase in client population for SWIM will make for an interesting process